EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L4	63	714/5.ccls. and @pd>="20061230"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:23
L5	35	714/11.ccls. and @pd>="20061230"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:23
L6	31	714/12.ccls. and @pd>="20061230"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:24
L7	33	714/13.ccls. and @pd>="20061230"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:24
L8	52	714/43.ccls. and @pd>="20061230"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:24
L9	0	cluster and ("neighbor's listing")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:29
L10	179	cluster and ("neighbor's list")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:29
L11	2	cluster and ("neighbor's list") and "virtual network"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:30
S1	29	(CHEN-MONSONG CHEN-MON-SONG MUKERJEE-BODHI CHEN-ALEX PAPPU-A PAPPU-APARNA).IN.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 08:38

EAST Search History

S2	8	("5805821" "5550577" "5712976" "6101547").PN.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 08:43
S3	1714	714/4.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 13:21
S4	353	714/4.ccls. and cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 13:21
S5	1022	714/5.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 13:42
S6	124	714/5.ccls. and cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 13:42
S7	675	714/11.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 14:54
S8	97	714/11.ccls. and cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 14:54
S9	292	714/12.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 15:02
S10	53	714/12.ccls. and cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 15:02
S11	547	714/13.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 15:11

EAST Search History

S12	106	714/13.ccls. and cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 15:11
S13	822	714/43.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 15:11
S14	70	714/43.ccls. and cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/30 15:11
S15	135530	cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 15:28
S16	8	hierarchial adj2 cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:35
S17	34	hierarchial$2 with cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:37
S18	0	("video on demand" "video ondemand") with cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:38
S19	1258	video with cluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:38
S20	0	S19 and "virtual multicast bus"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:39
S21	0	cluster and "virtual multicast bus"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:39

EAST Search History

S22	4	cluster and "virtual multicast"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:40
S23	37	cluster and "virtual bus"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:40
S24	8	cluster and "virtual bus" and multicast	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:42
S25	2178	cluster and ((node processor server computer) with priority)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 16:05
S26	597	S25 and ("failover" "fail over" "backup")	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/07/06 16:05
S27	79	S25 and ("failover" "fail over" "backup") and multicast	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 08:44
S28	60	S25 and ("failover" "fail over" "backup") and multicast and video	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 09:29
S29	2	"6728748".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/31 12:37

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L12	527	714/4.ccls.	US-PGPUB	OR	ON	2007/07/06 15:58
L13	1185	cluster and ((node processor server computer) with priority)	US-PGPUB	OR	ON	2007/07/06 16:05
L14	356	L13 and ("failover" "fail over" "backup")	US-PGPUB	OR	ON	2007/07/06 16:05

The "AND" operator is unnecessary -- we include all search terms by default. [details]

Web Results **1** - **10** of about **1,820,000** for **cluster AND (node OR processor OR server OR computer) AND**

Exchange **Server** 2003 **Cluster** Configuration Checklist
Clustering, by default, uses event log replication, which replicates all event occurrences from one **server node** in a **cluster** to all other nodes in the same ...
www.microsoft.com/technet/itshowcase/content/exch**cluster**cklist.mspx - 55k - Cached - Similar pages

Using Microsoft Virtual **Server** 2005 to Create and Configure a Two ...
Microsoft Virtual **Server** 2005 allows for two-**node clustering** of virtual ..... the Node2 for **computer** name for the **node** you want to add to the **cluster**. ...
www.microsoft.com/technet/prodtechnol/virtual**server**/deploy/cvs2005.mspx - 138k - Cached - Similar pages

Moab Administrator's Guide - **Node** Allocation Policies
The **node** allocation **priority** function can be specified on a **node** by **node** or **cluster** wide basis. In both cases, the recommended approach is to specify the ...
www.**cluster**resources.com/products/mwm/docs/5.2**node**allocation.shtml - 34k - Cached - Similar pages

Moab Administrator's Guide - Configuring **Node** Attributes
The default **node priority** is 0. A default **cluster**-wide **node priority** may be set by ... Knowing a **node's processor** speed can help the scheduler improve ...
www.**cluster**resources.com/products/mwm/docs/12.2**node**attributes.shtml - 35k - Cached - Similar pages

How to Configure **Cluster** Networking Components and **Priority**
This topic explains how to use **Cluster** Administrator to configure the **cluster** networking components and **priority** for a Microsoft Exchange **Server** 2007 ...
technet.microsoft.com/en-us/library/bb123936.aspx - 14k - Cached - Similar pages

Multiple **cluster** signal **processor** architecture - Patent 5392446
Each signal processing **cluster** comprises a system control **processor** ..... The system control **processor** 21 is a modular self-contained 32-bit **computer**. ...
www.freepatentsonline.com/5392446.html - 59k - Cached - Similar pages

Recommended private "Heartbeat" configuration on a **cluster server**
Communication between **Server Cluster** nodes is critical for smooth **cluster** operations. ... Sets the proper **Cluster** communication **priority** order. ...
support.microsoft.com/kb/258750 - Similar pages

Built For Speed: Develop Turbocharged Apps For Windows Compute ...
This article explores the services provided by Compute **Cluster Server** 2003 and the ... The IT department's support personnel control job **priority**, **node** ...
msdn.microsoft.com/msdnmag/issues/06/04/**Cluster**Computing/default.aspx - 84k - Cached - Similar pages

Modifying the /etc/lltab & a low **priority** link - **Cluster Server** ...

6) restart **cluster**. On my second point, should a low **priority** link be ... so pull the hi pri links and then change something on one **node** (add a new SG named ...

https://forums.symantec.com/syment/board/message?board.id=1&message.id=2418 - 61k - Cached - Similar pages

High-availability Service Management

Figure 1-11 shows an example of a high-availability **cluster** service that is a web **server** named "content-webserver". It is running in **cluster node** B and is **...**
www.redhat.com/docs/manuals/csgfs/browse/rh-cms-desc-ov-en/s1-service-management-overview.html - 8k - Cached - Similar pages

Google Groups results for **cluster AND (node OR processor OR server OR computer) AND priority**

Boost SQL **Server priority** on Windows switch - microsoft.public.sqlserver.clu ... - Jan 13, 2004
Routing interprocess connections between NLB **cluster ...** - microsoft.public.windows.serve ... - Jun 4, 2006
2007 A/P **cluster** networking? - microsoft.public.exchange.clus ... - Jun 22, 2007

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

Terms used:
luster AND node OR processor OR server OR computer AND priority

Found 88,642 of 205,978

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Distributed variable server for atomic unification

Alon Kleinman, Yoram Moses, Ehud Shapiro

August 1990 **Proceedings of the ninth annual ACM symposium on Principles of distributed computing PODC '90**

Publisher: ACM Press

Full text available: pdf(1.52 MB) Additional Information: full citation, references, citings, index terms

2 Exploiting perception in high-fidelity virtual environments: Exploiting perception in high-fidelity virtual environments

Additional presentations from the 24th course are available on the citation page

Mashhuda Glencross, Alan G. Chalmers, Ming C. Lin, Miguel A. Otaduy, Diego Gutierrez

July 2006 **ACM SIGGRAPH 2006 Courses SIGGRAPH '06**

Publisher: ACM Press

Full text available: pdf(5.07 MB) mov(68:6 MIN) Additional Information: full citation, appendices and supplements, abstract, references, cited by, index terms

The objective of this course is to provide an introduction to the issues that must be considered when building high-fidelity 3D engaging shared virtual environments. The principles of human perception guide important development of algorithms and techniques in collaboration, graphical, auditory, and haptic rendering. We aim to show how human perception is exploited to achieve realism in high fidelity environments within the constraints of available finite computational resources.In this course w ...

Keywords: collaborative environments, haptics, high-fidelity rendering, human-computer interaction, multi-user, networked applications, perception, virtual reality

3 Simulation potpouri: Simulation model of the cable data network for the analysis and evaluation of network performance

D. Gan, R. Paterson

December 1982 **Proceedings of the 14th conference on Winter Simulation - Volume 2 WSC '82**

Publisher: Winter Simulation Conference

Full text available: pdf(1.28 MB) Additional Information: full citation, abstract

A Cable Data Network (CDN) simulation model was developed on VAX 11/780 computer facility in PASCAL as a part of the MX-C³ system study. Its primary purpose was to supplement theoretical analysis and to evaluate the impact of changing the CDN (sub)system requirements on the performance measured primarily in terms of network reaction time and queue (buffer) buildup at the CDN nodes. The validated simulation model provided a powerful tool in rapidly determining the quantitati ...

4 A closed network with a discriminatory processor-sharing server

D. Mitra, A. Weiss

April 1989 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1989 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '89**, Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(792.94 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper gives a simple, accurate first order asymptotic analysis of the transient and steady state behavior of a network which is closed, not product-form and has multiple classes. One of the two nodes of the network is an infinite server and the discipline in the other node is discriminatory processor-sharing. Specifically, if there are nj jobs of class j at the latter node, then each class j job receives a fraction ...

5 Scheduling: Minimizing the stretch when scheduling flows of biological requests

Arnaud Legrand, Alan Su, Frédéric Vivien

July 2006 **Proceedings of the eighteenth annual ACM symposium on Parallelism in algorithms and architectures SPAA '06**

Publisher: ACM Press

Full text available: pdf(240.91 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we consider the problem of scheduling distributed biological sequence comparison applications. This problem lies in the divisible load framework with negligible communication costs. Thus far, very few results have been proposed in this model. We discuss and select relevant metrics for this framework: namely max-stretch and sumstretch. We explain the relationship between our model and the preemptive uni-processor case, and we show how to extend algorithms that have been proposed in ...

Keywords: competitive analysis, divisible load, flow time, linear programming, online algorithm, scheduling, stretch

6 Implicit coscheduling: coordinated scheduling with implicit information in distributed systems

Andrea Carol Arpaci-Dusseau

August 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 3

Publisher: ACM Press

Full text available: pdf(1.83 MB) Additional Information: full citation, abstract, references, citings, index terms

In modern distributed systems, coordinated time-sharing is required for communicating processes to leverage the performance of switch-based networks and low-overhead protocols. Coordinated time-sharing has traditionally been achieved with gang scheduling or explicit coscheduling, implementations of which often suffer from many deficiencies: multiple points of failure, high context-switch overheads, and poor interaction with client-server, interactive, and I/O -intensive workloads. I ...

Keywords: clusters, coscheduling, gang scheduling, networks of workstations, proportional-share scheduling, two-phase waiting

7 Piranha: a scalable architecture based on single-chip multiprocessing

Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture ISCA '00**, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

8 The V distributed system

David Cheriton

March 1988 **Communications of the ACM**, Volume 31 Issue 3

Publisher: ACM Press

Full text available: pdf(2.55 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The V distributed System was developed at Stanford University as part of a research project to explore issues in distributed systems. Aspects of the design suggest important directions for the design of future operating systems and communication systems.

9 LoPC: modeling contention in parallel algorithms

Matthew I. Frank, Anant Agarwal, Mary K. Vernon

June 1997 **ACM SIGPLAN Notices , Proceedings of the sixth ACM SIGPLAN symposium on Principles and practice of parallel programming PPOPP '97**, Volume 32 Issue 7

Publisher: ACM Press

Full text available: pdf(1.35 MB) Additional Information: full citation, abstract, references, citings, index terms

Parallel algorithm designers need computational models that take first order system costs into account, but are also simple enough to use in practice. This paper introduces the LoPC model, which is inspired by the LogP model but accounts for contention for message processing resources in parallel algorithms on a multiprocessor or network of workstations. LoPC takes the *L*, *o* and *P* parameters directly from the LogP model and uses them to predict the cost of contention, *C*

10 A survey of processors with explicit multithreading

Theo Ungerer, Borut Robič, Jurij Šilc

March 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 1

Publisher: ACM Press

Full text available: pdf(920.16 KB) Additional Information: full citation, abstract, references, citings, index terms

Hardware multithreading is becoming a generally applied technique in the next generation of microprocessors. Several multithreaded processors are announced by industry or already into production in the areas of high-performance microprocessors, media, and network processors.A multithreaded processor is able to pursue two or more threads of control in parallel within the processor pipeline. The contexts of two or more threads of control are often stored in separate on-chip register sets. Unused i ...

Keywords: Blocked multithreading, interleaved multithreading, simultaneous multithreading

11 A unifying framework for distributed simulation

R. Bagrodia, K. M. Chandy, Wen Toh Liao

October 1991 **ACM Transactions on Modeling and Computer Simulation (TOMACS)**, Volume 1 Issue 4

Publisher: ACM Press

Full text available: pdf(2.34 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A theory of distributed simulation applicable to both discrete-event and continuous simulation is presented. It derives many existing simulation algorithms from the theory and describes an implementation of a new algorithm derived from the theory. A high-level discrete-event simulation language has been implemented, using the new algorithm, on parallel computers; performance results of the implementation are also presented.

12 The Mercury Interconnect Architecture: a cost-effective infrastructure for high-performance servers

Wolf-Dietrich Weber, Stephen Gold, Pat Helland, Takeshi Shimizu, Thomas Wicki, Winfried Wilcke

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture ISCA '97**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(1.53 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents HAL's Mercury Interconnect Architecture, an interconnect infrastructure designed to link commodity microprocessors, memory, and I/O components into high-performance multiprocessing servers. Both shared-memory and message-passing systems, as well as hybrid systems are supported by the interconnect. The key attributes of the Mercury Interconnect Architecture are: low latency, high bandwidth, a modular and flexible design, reliability/availability/serviceability (RAS) features, ...

13 Process migration

Dejan S. Milojičić, Fred Douglis, Yves Paindaveine, Richard Wheeler, Songnian Zhou

September 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 3

Publisher: ACM Press

Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

Keywords: distributed operating systems, distributed systems, load distribution, process migration

14 Providing absolute differentiated services for real-time applications in static-priority scheduling networks

Shengquan Wang, Dong Xuan, Riccardo Bettati, Wei Zhao
April 2004 **IEEE/ACM Transactions on Networking (TON)**, Volume 12 Issue 2
Publisher: IEEE Press
Full text available: pdf(519.25 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we propose and analyze a methodology for providing absolute differentiated services for real-time applications. We develop a method that can be used to derive delay bounds without specific information on flow population. With this new method, we are able to successfully employ a utilization-based admission control approach for flow admission. This approach does not require explicit delay computation at admission time and, hence, is scalable to large systems. We assume the underlyi ...

Keywords: admission control, delay bound, differentiated services, priority assignment, real time, static-priority scheduling, utilization-based

15 Computer-aided parallelization of continuous media applications: the 4D beating heart slice server

J. Tárraga, V. Messerli, O. Figueiredo, B. Gennart, R. D. Hersch
October 1999 **Proceedings of the seventh ACM international conference on Multimedia (Part 1) MULTIMEDIA '99**
Publisher: ACM Press
Full text available: pdf(1.62 MB) Additional Information: full citation, abstract, references, index terms

Parallel servers for I/O and compute intensive continuous media applications are difficult to develop. A server application comprises many threads located in different address spaces as well as files striped over multiple disks located on different computers. The present contribution describes the construction of a continuous media server, the 4D beating heart slice server, based on a computer-aided parallelization tool (CAP) and on a library of parallel file system components enabling the ...

Keywords: 4D tomographic images, computer-aided parallelization, disk scheduling, parallel I/O streaming, parallel continuous media server, resource reservation

16 Shape-based retrieval and analysis of 3D models

Thomas Funkhouser, Michael Kazhdan
August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**
Publisher: ACM Press
Full text available: pdf(12.56 MB) Additional Information: full citation, abstract

Large repositories of 3D data are rapidly becoming available in several fields, including mechanical CAD, molecular biology, and computer graphics. As the number of 3D models grows, there is an increasing need for computer algorithms to help people find the interesting ones and discover relationships between them. Unfortunately, traditional text-based search techniques are not always effective for 3D models, especially when queries are geometric in nature (e.g., find me objects that fit into thi ...

17 Scalable concurrent priority queue algorithms

Nir Shavit, Asaph Zemach
May 1999 **Proceedings of the eighteenth annual ACM symposium on Principles of distributed computing PODC '99**
Publisher: ACM Press
Full text available: pdf(1.35 MB) Additional Information: full citation, references, citings, index terms

18 Finite-grain message passing concurrent computers

W. J. Dally

January 1988 **Proceedings of the third conference on Hypercube concurrent computers and applications: Architecture, software, computer systems, and general issues - Volume 1**

Publisher: ACM Press

Full text available: pdf(1.22 MB) Additional Information: full citation, abstract, references, citings, index terms

Fine-grain concurrent computers, by operating at a fine grain, increase the amount of concurrency that can be efficiently exploited in a given problem. Programming is simplified because programs may be partitioned into natural units of methods and objects and these objects are addressed uniformly whether they are local or remote. The construction of these machines poses challenging problems in reducing overhead, increasing communication bandwidth, and developing resource management techniqu ...

19 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research CASCON '97**

Publisher: IBM Press

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

20 A comparative study of parallel and sequential priority queue algorithms

Robert Rönngren, Rassul Ayani

April 1997 **ACM Transactions on Modeling and Computer Simulation (TOMACS)**, Volume 7 Issue 2

Publisher: ACM Press

Full text available: pdf(640.10 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Priority queues are used in many applications including real-time systems, operating systems, and simulations. Their implementation may have a profound effect on the performance of such applications. In this article, we study the performance of well-known sequential priority queue implementations and the recently proposed parallel access priority queues. To accurately assess the performance of a priority queue, the performance measurement methodology must be appropriate. We use the Classic ...

Keywords: parallel access priority queue, pending event set implementations, priority queue